

# Analyzing Tag Distributions in Folksonomies for Resource Classification

Arkaitz Zubiaga, Raquel Martínez, and Víctor Fresno

NLP & IR Group @ UNED

**Abstract.** Recent research has shown the usefulness of social tags as a data source to feed resource classification. Little is known about the effect of settings on folksonomies created on social tagging systems. In this work, we consider the settings of social tagging systems to further understand tag distributions in folksonomies. We analyze in depth the tag distributions on three large-scale social tagging datasets, and analyze the effect on a resource classification task. To this end, we study the appropriateness of applying weighting schemes based on the well-known TF-IDF for resource classification. We show the great importance of settings as to altering tag distributions. Among those settings, tag suggestions produce very different folksonomies, which condition the success of the employed weighting schemes. Our findings and analyses are relevant for researchers studying tag-based resource classification, user behavior in social networks, the structure of folksonomies and tag distributions, as well as for developers of social tagging systems in search of an appropriate setting.

## 1 Introduction

Social tagging systems allow users to annotate resources with tags in an aggregated way. These systems generate large amounts of metadata that increase the availability of resource descriptions, and can enhance information access in a wide sense [9, 3]. It has attracted a large number of researchers to using them for improving resource retrieval, organization and classification tasks, among others [8]. Regarding resource classification tasks, social tags have shown high effectiveness by outperforming content-based approaches [16, 13, 6]. Social tags represent an interesting data source for the sake of resource classification tasks.

However, each social tagging system may produce a different folksonomy structure. So far, research on tag-based classification has focused on analyzing the use of social tags for specific datasets, and no attention has been paid to settings of different systems [16, 6]. We aim at complementing earlier research by further analyzing tag distributions as a way of weighting the relevance of tags in the collection. As an approach to considering distributions as an indicator of the relevance of tags in a collection, we perform a classification study by adapting the TF-IDF weighting scheme [14]. Text document collections are simply made up by documents containing terms, whereas social tagging systems involve different users and bookmarks for each resource containing tags. Several

works have applied the TF-IDF scheme to social tagging systems [2, 15, 11, 12], but there is no evidence on its suitability for these systems.

In this paper, we analyze the settings, and study folksonomies and tag distributions on three large-scale social tagging datasets. We apply three weighting schemes based on TF-IDF to define the relevance of tags on a folksonomy, and explore their suitability according to system’s settings. We evaluate the results by performing resource classification experiments, comparing those weighting schemes to the sole use of TF in the absence of weights. We find that tag distributions in folksonomies can help determine relevance of tags, but they are strictly subject to settings of the tagging system. Specifically, an IDF-like weighting scheme does not present the desired effect when a social tagging system suggests tags to the user. where utterly different folksonomy structures are produced.

The paper is organized as follows. Next, we define a collaborative tagging system in Section 2. We detail the studied collaborative tagging systems, the process of generation of the datasets, and perform a thorough analysis of tag distributions in Section 3. We provide a brief overview of TF-IDF, and introduce the variants adapted to collaborative tagging in Section 4. Then, we present the tag-based resource classification experiments and analyze the results in Section 5. We conclude the paper with our thoughts in Section 6.

## 2 Social Tagging

Tagging is an open way that allows users to bookmark and annotate with tags their favorite resources (e.g., web pages, movies or books). Tagging resources facilitates future retrieval by relying on tags as metadata describing resources. On social tagging systems, users can collaboratively annotate resources, so that many users can tag the same resource. Tags by different users aggregated on a resource provide vast amounts of metadata. The collection of tags defined by users of a system creates a tag-based organization, so-called folksonomy. The subset of tags assigned by a single user creates a smaller folksonomy, also known as personomy. For instance, CiteULike<sup>1</sup>, LibraryThing and Delicious are social tagging system where users collaboratively annotate resource. On these systems, each resource (papers, books and URLs, respectively) can be annotated and tagged by all the users who consider it interesting. On a social tagging system, there is a set of users ( $U$ ), who are posting bookmarks ( $B$ ) for resources ( $R$ ) annotated by tags ( $T$ ). Each user  $u_i \in U$  can post a bookmark  $b_{ij} \in B$  of a resource  $r_j \in R$  with a set of tags  $T_{ij} = \{t_{ij1}, \dots, t_{ijp}\}$ , with a variable number  $p$  of tags. After  $k$  users posted  $r_j$ , it is described with a weighted set of tags  $T_j = \{w_{j1}t_{j1}, \dots, w_{jn}t_{jn}\}$ , where  $w_{j1}, \dots, w_{jn} \leq k$  represent the number of assignments of a specific tag. Accordingly, each bookmark is a triple of a user, a resource, and a set of tags:  $b_{ij} : u_i \times r_j \times T_{ij}$ . Thus, each user saves bookmarks of different resources, and a resource has bookmarks posted by different users. The result of aggregating tags within bookmarks by a user is known as the personomy of the

<sup>1</sup> <http://www.citeulike.org>

user:  $T_i = \{w_{i1}t_{i1}, \dots, w_{im}t_{im}\}$ , where  $m$  is the number of different tags in user's personomy.

For instance, a user could tag this work as **social-tagging**, **research**, and **paper**, whereas another user could use the tags **classification**, **paper**, **social-tagging**, and **social-bookmarking** to annotate it. The aggregation of these annotations would produce the following: **social-tagging** (2), **paper** (2), **social-bookmarking** (1), **classification** (1), and **research** (1). In this example the values represent the weighted union of all tags.

### 3 Datasets

Next, we introduce and analyze the three large-scale datasets we gathered from well-known social tagging sites: Delicious, LibraryThing, and GoodReads. All of them had been gathered from March to May 2010.

#### 3.1 Studied Social Tagging Sites

**Delicious** is a social tagging site that allows users to save and tag their preferred web pages, in order to ease the subsequent navigation and retrieval on large collections of annotated bookmarks. On a social bookmarking site, any web page can be saved, so that the range of covered topics can become as wide as the Web is. It is known that the site is biased to some computer and design related topics though. When a user saves a URL as a bookmark, the system suggests tags previously used for that URL if some users had annotated it before. Thus, new annotators can easily add tags used by earlier users without typing them.

**LibraryThing** and **GoodReads** are social tagging sites where users save and annotate books. Commonly, users annotate the books they own, they have read, or they are planning to read. Besides readers, there are also well-known writers and libraries contributing as users on these sites. The main difference among these two systems is that LibraryThing does not suggest tags when saving a book, whereas GoodReads lets the user select from tags within his personomy, that is, tags he previously assigned to other books. The latter makes it easier to reuse users' favorite tags, without re-typing them. Another difference is that LibraryThing allows some users to group tags with the same meaning, linking thus typos, synonyms and translations to a single tag, e.g., *science-fiction*, *sf* and *ciencia ficción* are grouped into *science fiction*.

Despite of the aforementioned differences, all of them have some characteristics in common: users save resources as bookmarks, a bookmark can be annotated by a variable number of tags ranging from zero to unlimited, and the vocabulary of the tags is open. Table 1 summarizes the features of the social tagging sites we study in this work.

	<b>Delicious</b>	<b>LibraryThing</b>	<b>GoodReads</b>
<b>Resources</b>	web documents	books	books
<b>Tag suggestions</b>	tags from earlier bookmarks of the resource	no	tags in user's personomy
<b>Users</b>	general	readers, writers & libraries	readers, writers & libraries
<b>Tag grouping</b>	no	selected users suggest merging tags	no
<b>Vocabulary</b>	open	open	open
<b>Tag insertion</b>	space-separated	comma-separated	one by one text-box

**Table 1.** Characteristics of studied social tagging systems.

### 3.2 Generation Process of Datasets

First of all, we queried the three sites for popular resources<sup>2</sup>. In the case of Delicious, we gathered a set of 87,096 URLs. As regards to LibraryThing and GoodReads, we found an intersection of 65,929 popular books. In the next step, we looked for classification labels assigned by experts for both kind of resources. For Delicious, we used the Open Directory Project<sup>3</sup> (ODP) as a ground truth, where we found an intersection of 12,616 URLs with our set. For the set of books, we fetched their classification for both the Dewey Decimal Classification (DDC) and the Library of Congress Classification (LCC) systems. The former is a classical taxonomy that is still widely used in libraries, whereas the latter is used by most research and academic libraries. We found that 27,299 books were categorized on DDC, and 24,861 books have an LCC category assigned.

Finally, we queried (a) Delicious for all the public bookmarks from users involved in the set of categorized URLs, and (b) LibraryThing and GoodReads for all the public bookmarks from users involved in the set of categorized books. This resulted in a large collections of bookmarks, not only for categorized resources, but also for many others. Each bookmark comprises: (1) the user who annotated, (2) the annotated resource, and (3) the associated tags. We saved all the tags attached to each bookmark, except for GoodReads, where a tag is automatically attached to each bookmark depending on the reading state of the book: *read*, *currently-reading* or *to-read*. We do not consider this to be part of the tagging process, but just an automated step, and we removed all their appearances in the dataset.

<sup>2</sup> We consider a resource to be popular if at least 100 users have annotated it as a bookmark. It was shown that the tag set of a resource tends to converge when that many users contribute to it [7].

<sup>3</sup> <http://www.dmoz.org>

### 3.3 Statistics and Analysis of the Datasets

It is worthwhile noting that attaching tags to a bookmark is an optional step, so that depending on the social tagging site, a number of bookmarks may remain without tags. Table 2 presents the number of users, bookmarks and resources we gathered for each of the datasets, as well as the percent with attached annotations. In this work, as we rely on tagging data, we only consider annotated data, ruling out bookmarks without tags. Thus, from now on, all the results and statistics presented are based on annotated bookmarks. From these statistics, it stands out that most users (above 87%) provide tags for bookmarks on Delicious, whereas there are fewer users who tend to assign tags to resources on LibraryThing and GoodReads (roughly 38% and 17%, respectively).


<b>Delicious</b>			
	Annotated	Total	Percent
Users	1,618,635	1,855,792	87.22%
Bookmarks	273,478,137	300,571,231	91.00%
Resources	92,432,071	102,828,761	89.89%
Tags		11,541,977	-
<b>LibraryThing</b>			
	Annotated	Total	Percent
Users	153,606	400,336	38.37%
Bookmarks	22,343,427	44,612,784	50.08%
Resources	3,776,320	5,002,790	75.48%
Tags		2,140,734	-
<b>GoodReads</b>			
	Annotated	Total	Percent
Users	110,344	649,689	16.98%
Bookmarks	9,323,539	47,302,861	19.71%
Resources	1,101,067	1,890,443	58.24%
Tags		179,429	-

**Table 2.** Statistics on availability of tags in users, bookmarks, and resources for the three datasets.

Regarding the distribution of tags across all the resources, users and bookmarks in the datasets, there is a clear difference of behavior among the three collections. Figure 1 shows the percent of resources, users and bookmarks on which tags are annotated according to their rank on the system. That is, the X axis refers to the percent of the tag rank, whereas the Y axis represents the percent of appearances in resources, users and bookmarks. For instance, if the tag ranked first had been annotated on the half of the resources, the value for the top ranked tag on resources would be 50%. Thus, these graphs enable to analyze how popular are the tags in the top as compared to the tags in the tail on each site. It stands out that GoodReads has the highest usage of tags in the

tail, but Delicious presents the highest usage of tags in the top. Delicious is the site with highest diversity of tags, where a few tags become really popular, and many tags are seldom-used (note the logarithmic scale of the graphs). We believe that the reasons for these differences on tag distributions are:

- Since Delicious suggests tags that have been annotated by previous users to a resource, it is obvious that those tags on the top are likely to happen more frequently, whereas others may barely be used.
- LibraryThing and GoodReads do not suggest tags used by earlier users and, therefore, tags other than those on the top tend to be used more frequently than in Delicious.
- GoodReads suggests tags from previous bookmarks of the same user, instead of tags that others assigned to the resource being tagged. Thus, this encourages reusing tags in their personomy, making it remain with a smaller number of tags (see Table 3). In addition, users tend to assign fewer tags to a bookmark on average, probably due to the one-by-one tag insertion of site’s interface.



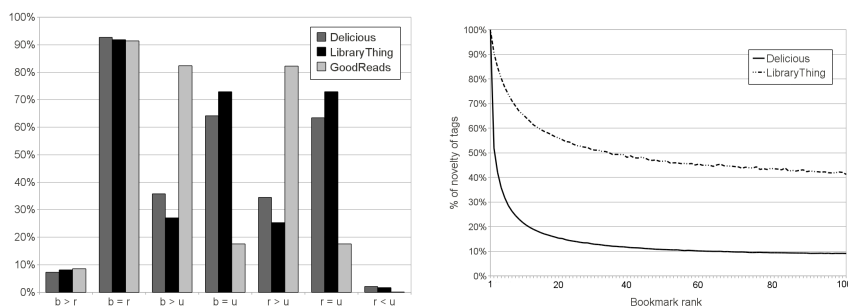
**Fig. 1.** Tag usage percentages. These 3 graphs represent, on a logarithmic scale for both  $x$  and  $y$  axes, the percent of annotations to resources, users, and bookmarks per tag rank on each site.

# of tags	Delicious	LibraryThing	GoodReads
Per resource	33.35	14.53	13.33
Per user	632.714	357.15	131.03
Per bookmark	3.75	2.46	1.55

**Table 3.** Average counts of different tags

Regarding the distribution of tags across resources, users, and bookmarks, Figure 2(a) shows percents of tags appearing more, equal or less frequently in

an item (i.e., resources, users or bookmarks) than in another. It is obvious that a tag cannot appear in smaller number of bookmarks than users or resources, by definition. Looking at the rest of data, it stands out that tags tend to appear in more bookmarks than users ( $b > u$ ) and more resources than users ( $r > u$ ) for GoodReads, due to the same feature that allows users to select among tags in their personomy. However, LibraryThing and Delicious have many tags present in the same number of bookmarks and users ( $b = u$ ), and resources and users ( $r = u$ ), even though the difference is more marked for the former site. This reflects the large number of tags that users utilize just once on these sites. All three sites have two features in common: there are a few exceptions of tags utilized by more users than the number of resources it appears in ( $r < u$ ), and almost all the tags are present in the same number of bookmarks and resources ( $b = r$ ). The latter, combined with the lower ( $b = u$ ) values, means that there is a large number of users spreading personal tags across resources that only have a bookmark with that tag, especially on GoodReads, but also for the other two sites.



(a) Tag distribution across resources ( $r$ ), (b) Novelty ratio of tags per rank of users ( $u$ ) and bookmarks ( $b$ ).

Finally, we analyze the evolution of annotations throughout the time, i.e., the extent to which a bookmark introduces new tags into a resource that were not present in earlier bookmarks. Figure 2(b) shows these statistics for Delicious and LibraryThing. The plot for GoodReads is not shown because neither the timestamp nor the ordering of the bookmarks is available in our dataset. The graph shows, on average, the ratio of new tags, not present in earlier bookmarks of a resource, assigned in bookmarks that rank from first to 100th bookmark, i.e., if  $tag_1$  and  $tag_2$  were annotated in the first bookmark of a resource, and  $tag_2$  and  $tag_3$  in the second bookmark for the same resource, the ratio of novelty for the second bookmark is of 50%. It stands out the marked inferiority of tag novelty on Delicious as against to LibraryThing. This is, again, due to the tag suggestion policy of Delicious, what makes previous tags to re-occur more frequently.

## 4 Tag Weighting Functions

### 4.1 TF-IDF as a Term Weighting Function

TF-IDF is a term weighting function that combines the term frequency and inverse document frequency to produce a composite weight for each term in a document [14]. This weight is higher when  $t$  occurs many times within a small number of documents (thus contributing to high discriminating power to those documents); lower when the term occurs fewer times in a document, or occurs in many documents (thus offering a less pronounced relevance signal); and it becomes null when the term occurs in all the documents.

TF-IDF is the most widely used weighting function for representing documents in a vector-space model, and has become a “de facto standard” for automated classification tasks [10].

### 4.2 TF-IDF for Tags

In order to analyze the effect of tag distributions on each social tagging system to the resource classification task, we adapt the classical TF-IDF function. Unlike classical collections of web documents or library catalogs, where the terms distribute across documents in the collection, social tagging systems provide further data as new dimensions to explore into. Besides tags’ distribution across documents or annotated resources, different users set those tags within different bookmarks. These two characteristics are specific to social tagging, and were not available on classical text collections.

Next, we introduce three tag weighting approaches, taking the classical TF-IDF approach to the social tagging scenario, and adapting it to rely on resources, users and bookmarks. These three dimensions suggest defining that many tag weighting functions considering inverse resource frequency (IRF), inverse user frequency (IUF), and inverse bookmark frequency (IBF) values, respectively. The three approaches we compare in this work follow the same function for the tag  $i$  within the resource  $j$ :

$$TF\text{-}IxF_{ij} = tf_{ij} \cdot ixf$$

where  $tf_{ij}$  is the number of occurrences of the tag  $i$  in the resource  $j$ , and  $ixf$  is the inverse frequency function considered in each case,  $irf$ ,  $iuf$  or  $ibf$ , thus  $x$  being  $r$ ,  $u$ , or  $b$ . Accordingly, we define TF-IRF for resources, TF-IUF for users, and TF-IBF for bookmarks.

There is little work using TF-IRF, usually referred to as TF-IDF. [2] rely on this measure to infer similarity of tags by creating a tag graph, weighting the TF-IDF value of each user to a tag. [15] and [12] use this measure to represent the resources in a recommendation system where resources are recommended to users. [11] create tag vectors using TF-IDF to compute the similarity between two documents annotated on Delicious.

TF-IUF was inferred from a previous application to a collaborative filtering system [4]. With the aim of recommending resources to users, [5] and [12] rely



on the IUF for discovering similarities among users. The latter use both IUF and IRF to represent users and resources, respectively, but no comparison is performed among their characteristics. In [1], TF-IUF is used along with TF-IRF over Flickr for finding landmark photos.

To the best of our knowledge, TF-IBF has never been used so far. Even though all three frequencies can somehow be related, there are substantial differences among them. A tag used by many users can spread across many resources, or it can just congregate in a few resources. Likewise, this factor might affect the number of bookmarks.

## 5 Tag-based Classification

Next, we present the classification experiments that enable (1) the analysis of how each of the tag weighting functions contributes to the classification of annotated resources and whether they outperform the baseline relying only on the tag frequency (TF), as well as (2) discovering whether the settings of social tagging systems affect the performance of the tag weighting functions.

In order to analyze the impact of weighting functions on tag-based classification, we use multiclass Support Vector Machines (SVM) [10]. Specifically, we use "svm-multiclass"<sup>4</sup>, with the linear kernel and the default parameters.

We perform the classification tasks relying on the top level of the taxonomies. We maintained the structure of all the taxonomies, but merged the categories E (*History of America*) and F (*History of the United States and British, Dutch, French, and Latin America*) on LCC, as the differences between them do not seem clear. Thus, ODP is composed by 17 categories, DDC by 10, and LCC by 20. We used different training set sizes, and made 6 different selections for each size, getting the average of all 6 runs not to make the results depend on the selected training instances.

Table 4 shows the results of using tag weighting functions. For Delicious, the outperformance of TF shows that the use of weighting functions is not useful in this case. Going further into the analysis of the performance of representations relying on weighting functions, the results show that IUF gets the best results among them, followed by IBF, and then IRF. We believe that resource-based tag suggestions, as occurs on Delicious, are not helpful to this end. It makes the top tags become even more popular and it alters the natural distribution of tags. Thus, such a forced distribution of tags produces weights that score lower performances. Moreover, the fact that IUF is the best weighting function in this case, shows the importance of users who make their own choices instead of relying on suggestions. That is, users who differ from suggestion-based annotations give rise to higher weights for their seldom tags, which performs better than IRF and IBF.

In the case of LibraryThing, all the weighting functions are clearly superior to TF, since the former always outperform the latter. This shows that the studied

<sup>4</sup> [http://svmlight.joachims.org/svm\\_multiclass.html](http://svmlight.joachims.org/svm_multiclass.html)

<b>Delicious - ODP</b>							
training set	600	1400	2200	3000	4000	5000	6000
TF	<b>.533</b>	<b>.600</b>	<b>.629</b>	<b>.647</b>	<b>.660</b>	<b>.669</b>	<b>.680</b>
TF-IRF	.516	.571	.593	.607	.619	.631	.639
TF-IBF	.519	.573	.596	.611	.622	.633	.641
TF-IUF	.528	.580	.607	.625	.636	.653	.661
<b>LibraryThing - DDC</b>							
training set	3000	6000	9000	12000	15000	18000	21000
TF	.861	.864	.864	.867	.869	.869	.868
TF-IRF	.877	.889	.894	<b>.897</b>	<b>.900</b>	.902	.902
TF-IBF	.877	.889	.894	<b>.897</b>	<b>.900</b>	<b>.903</b>	<b>.904</b>
TF-IUF	<b>.881</b>	<b>.891</b>	<b>.895</b>	<b>.897</b>	.899	.901	.900
<b>LibraryThing - LCC</b>							
training set	3000	6000	9000	12000	15000	18000	21000
TF	.853	.857	.856	.861	.861	.857	.861
TF-IRF	.867	<b>.883</b>	.887	<b>.893</b>	.895	.894	.897
TF-IBF	.867	<b>.883</b>	<b>.888</b>	<b>.893</b>	<b>.896</b>	<b>.895</b>	<b>.898</b>
TF-IUF	<b>.871</b>	.882	.885	.892	.893	.892	.894
<b>GoodReads - DDC</b>							
training set	3000	6000	9000	12000	15000	18000	21000
TF	.745	.747	.754	.757	.757	.757	.756
TF-IRF	<b>.800</b>	.808	.813	<b>.817</b>	.816	.817	.816
TF-IBF	<b>.800</b>	<b>.809</b>	<b>.814</b>	<b>.817</b>	<b>.817</b>	<b>.818</b>	<b>.818</b>
TF-IUF	.797	.805	.810	.814	.814	.814	.814
<b>GoodReads - LCC</b>							
training set	3000	6000	9000	12000	15000	18000	21000
TF	.725	.731	.737	.738	.734	.731	.743
TF-IRF	<b>.781</b>	<b>.792</b>	<b>.797</b>	.801	<b>.802</b>	.799	.804
TF-IBF	<b>.781</b>	<b>.792</b>	<b>.797</b>	<b>.803</b>	<b>.802</b>	<b>.800</b>	<b>.805</b>
TF-IUF	.776	.788	.792	.797	.797	.794	.800

**Table 4.** Accuracy results of classification using weighting functions.

inverse weighting functions can be really useful for folksonomies created in the absence of suggestions. Tag weighting functions have successfully set suitable weights towards a definition of the representativity of tags in this case, in contrast to Delicious. Among the tag weighting functions, all of them perform similarly, and no clear outperformances can be seen in these results. However, IBF seems to provide slightly better results than the other two approaches, followed by IRF. IUF is the worst function in this case, suggesting that the number of users choosing each tag is not the most relevant feature when no suggestions are given.

On GoodReads, tag weighting functions also clearly outperform the sole use of TF. As on LibraryThing, IBF performs the best among the weighting functions, followed by IRF, and then IUF. Even though there are suggestions on GoodReads too, they rely on users' tags, and thereby these suggestions can only be applied to different resources. This shows that the effect of personomy-based suggestions is much smaller, and it affects to a lower extent or does not almost affect the distribution of tags, because suggestions do not spread to users.

Summarizing, results show that the studied inverse tag weighting functions can be really useful for determining the representativity of tags within the collection. However, folksonomies can suffer from resource-based tag suggestions, transforming the structure and distributions of folksonomies. This transformation can even be harmful for the definition of the tag weighting functions, and can bring about worse performance results than simply relying on TF, as happened on Delicious. Otherwise, in the absence of resource-based tag suggestions, the use of tag weighting functions improves the performance.

Comparing the results scored by tag weighting functions, it can be seen that IBF is always slightly better than IRF. The former is more detailed than the latter, because it considers the exact number of appearances of the tag besides the number of resources it appears in. Actually, IBF is the best approach for both LibraryThing and GoodReads, where there are no suggestions, or suggestions rely on user's personomy. When these suggestions rely on tags annotated earlier to the resource, as on Delicious, IUF performs better than the other two weighting functions, showing the relevance of users' ability to dismiss suggestions. However, even IUF is unable to outperform TF in this case.

## 6 Conclusions

This work complements earlier research on tag-based resource classification by exploring the effect of tag distributions on systems with different settings. Adapting a weighting scheme such as TF-IDF helps better understand these distributions. To the best of our knowledge, this is the first research work exploiting tags from different systems for a resource classification task. The study on three large-scale datasets from systems with different settings has given rise to better understanding the classification performance on each system. We have performed a thorough analysis of tag distributions on the datasets, and have analyzed the effect of those distributions as to the classification of resources using tags.

Among the settings, tag suggestions have shown to influence the structure of folksonomies greatly. Users tend to choose among suggested tags rather than providing their own tags, producing different folksonomies where a few tags stand out as compared to the rest. Systems with resource-based suggestions produce different tag distributions with respect to the absence of suggestions. The use of an IDF-like scheme shows that it cannot help determine the relevance of tags in such cases, where TF yields better results. However, we have found that tag weighting functions clearly outperform the TF approach when resource-based tag suggestions are disabled, i.e., on LibraryThing and GoodReads. Our findings are relevant for scientists studying tag-based resource classification, user behavior in social networks, the structure of folksonomies and tag distributions, as well as for developers of tagging systems in search of the appropriate setting.

Future work includes studying the suitability of inverse tag weighting functions for other information management tasks, and further analyzing suggestion biases like Delicious, in search of a weighting function to fit them.

## References

1. Abbasi, R., Chernov, S., Nejdl, W., Paiu, R., Staab, S.: Exploiting flickr tags and groups for finding landmark photos. In: ECIR '09: 31st European Conference on IR Research on Advances in Information Retrieval. pp. 654–661. Springer-Verlag, Berlin, Heidelberg (2009)
2. Angelova, R., Lipczak, M., Milios, E., Pralat, P.: Characterizing a social bookmarking and tagging network. In: ECAI 2008 Workshop on Mining Social Data (MSoDa). pp. 21–25. IOS (2008)
3. Awawdeh, R., Anderson, T.: Improving search in tag-based systems with automatically extracted keywords. In: Proceedings of KSEM '10, the 4th International Conference on Knowledge Science, Engineering and Management. pp. 378–387. KSEM'10, Springer-Verlag, Berlin, Heidelberg (2010)
4. Breese, J.S., Heckerman, D., Kadie, C.: Empirical analysis of predictive algorithms for collaborative filtering. In: 14th Conference on Uncertainty in Artificial Intelligence. pp. 43–52. Morgan Kaufmann (1998)
5. Diederich, J., Iofciu, T.: Finding communities of practice from user profiles based on folksonomies. In: 1st International Workshop on Building Technology Enhanced Learning solutions for Communities of Practice (2006)
6. Godoy, D., Amandi, A.: Exploiting the social capital of folksonomies for web page classification. In: Software Services for E-World. IFIP Advances in Information and Communication Technology, vol. 341, pp. 151–160. Springer (2010)
7. Golder, S., Huberman, B.A.: The structure of collaborative tagging systems. *Journal of Information Science*, 32(2) pp. 198–208 (2006)
8. Gupta, M., Li, R., Yin, Z., Han, J.: Survey on social tagging techniques. *SIGKDD Explorations* 12(1), 58–72 (2010)
9. Heymann, P., Koutrika, G., Garcia-Molina, H.: Can social bookmarking improve web search? In: WSDM '08, the international conference on Web search and web data mining. pp. 195–206. ACM, New York, NY, USA (2008)
10. Joachims, T.: Text categorization with support vector machines: Learning with many relevant features. In: ECML. pp. 137–142. Springer, Berlin (1998)
11. Li, X., Guo, L., Zhao, Y.E.: Tag-based social interest discovery. In: WWW '08: 17th international conference on World Wide Web. pp. 675–684. ACM, New York, NY, USA (2008)
12. Liang, H., Xu, Y., Li, Y., Nayak, R., Tao, X.: Connecting users and items with weighted tags for personalized item recommendations. In: HT '10: 21st ACM conference on Hypertext and hypermedia. pp. 51–60. ACM, New York, NY, USA (2010)
13. Noll, M.G., Meinel, C.: The metadata triumvirate: Social annotations, anchor texts and search queries. In: Web Intelligence and Intelligent Agent Technology, 2008. WI-IAT '08. IEEE/WIC/ACM International Conference on. vol. 1, pp. 640–647 (2008)
14. Salton, G., Buckley, C.: Term-weighting approaches in automatic text retrieval. *Information Processing and Management* 24, 513–523 (August 1988)
15. Shepitsen, A., Gemmell, J., Mobasher, B., Burke, R.: Personalized recommendation in social tagging systems using hierarchical clustering. In: RecSys '08. pp. 259–266. ACM, New York, NY, USA (2008)
16. Zubiaga, A., Martínez, R., Fresno, V.: Getting the most out of social annotations for web page classification. In: DocEng '09: 9th ACM symposium on Document engineering. pp. 74–83. ACM, New York, NY, USA (2009)